

REMARKS

Claims 1-8 were presented for examination and were pending in this application. In the latest Office Action, claims 1-8 were rejected. On the basis of the following remarks, consideration of this application and allowance of all pending claims are requested.

Claims 1 and 3-8 were rejected as anticipated by U.S. Patent No. 6,556,548 to Kirkby et al. Claim 2 was rejected as made obvious by Kirkby in view of U.S. Patent No. 6,246,692 to Dai et al. Applicant respectfully traverses the rejections.

The claims recite a network system comprising a set of nodes, with variable capacity connections transporting data from the source nodes to the destination nodes of the network, wherein the destination nodes of the connections control the capacities of the connections according to the data traffic loads of the connections. Kirkby does not disclose (or suggest, in combination with Dai) the claimed network system for a number of reasons, at least some of which are provided below.

“the capacity of each connection [is] controlled from its destination node”

Claim 1 recites that the capacity of each connection configured to transport data from the source node to the destination node is controlled by its destination node. Kirkby does not disclose a network system in which the connection capacities are controlled from their destination nodes.

For instance, in Kirkby, column 5, lines 10-18, state that “It is the network manager (central controller) or the local network element manager (distributed controller) . . . which calculates each user’s allocated bandwidth.” In column 15, lines 63-64, Kirkby further states, “The local controllers divide the users’ WtP values according to the relative demands of the

resources en route.” As this passage in Kirkby makes clear, the network bandwidth allocation involves at least either a *central* controller or the *local* (i.e., *source*) nodes. Unlike the system recited in claim 1, Kirkby’s *destination* nodes do not control the connection capacities.

Moreover, Kirkby states in column 7, lines 1-8, that “demand for a resource . . . is taken over all the source-sinks using the resource,” and in column 8, lines 5-6, that “link (resource) 1-10 carries routes 1-10, 1-10-18.” In these passages, Kirkby explains that a given network link (i.e., a connection resource) can be allocated to connections to *at least two different* destination nodes (sinks) (e.g., the destination nodes 10 and 18, in that case of the previous quotation from Kirkby). In fact, it would not be technically feasible for the destination nodes to control the connection capacities where a given network resource can be used by or allocated to connections to *multiple* different destination nodes. Such a configuration would lead to resource allocation conflicts if the individual destination nodes were to control connection capacities independently while the same network capacity resources could be allocated to connections with *different* destination nodes.

Furthermore, Kirkby explains at column 14, lines 3-4, “Each controller needs to know the relative demands on the network resources.” Kirkby further explains at column 16, lines 11-23:

At a particular node, there is through traffic . . . as well as ingress traffic . . . In order to calculate the demand at a resource, the total price paid to that resource by all the flows which are using that resource is required. There will be two proportional price paid parameters to a resource, one from the proportional price paid by the ingress flow . . . the other from the proportional price paid by the through traffic . . . The demand at a resource is the sum of the above two proportional prices paid to that resource.

Kirkby’s controller thus needs information about multiple network resources and their local *through* traffic and *ingress* traffic to allocate network resources at a particular node. In contrast, each destination node recited in claim 1 needs to know the traffic loads of only those connections for which it is a destination node (i.e., the connections controlled by that destination node).

Unlike in Kirkby, for the purpose of connection capacity control, the claimed nodes do not need to know anything about *through* or *ingress* traffic at that node or about network resources other than its own network *egress* capacity.

For at least the above reasons Kirkby does not disclose the network system of claim 1, where the connection capacities are controlled from their individual destination nodes. This difference provides significant advantages over Kirkby, including reduced network complexity due to the elimination of the need to have multiple network nodes or a centralized controller to participate in the process of controlling capacity of a given connection. This results in improved scalability, predictability and reliability of the network, as well as cost-efficiency and performance benefits.

“based at least in part on the traffic loads associated with the connections”

Claim 1 further recites that the destination node’s control of each connection is “based at least in part on the traffic loads associated with the connections configured to transport data to that destination node.” Kirkby does not disclose this feature.

In the Office Action, the examiner cited the network element managers in Kirkby’s FIG. 4 that pass relative demand information to regulate resource allocation. In so doing, the examiner has interpreted Kirkby’s “demand information” to the claimed traffic loads. Applicant respectfully points out that this is an incorrect interpretation of Kirkby. In column 3, lines 38-45, Kirkby explains:

In this method, users express their willingnesses to pay [WtP] for network resources. A controller determines how the users' willingnesses to pay are to be divided between the resources in order to determine the relative demands for the resources. Each resource is then divided between those users using it in proportion to how much they are willing to pay for the use of their share of it.

In column 6, lines 62-66, Kirkby further states, “The method iterates from these starting points to establish a better allocation of resources and price (demand) between the competing users”, and in column 13, lines 41-44, states, “The controller at each node uses the available information about the demands, in order to calculate the proportional division of the price paid to (demand for) each resource.” Accordingly, Kirkby uses the term “demand” to refer to the *price a user is willing to pay* for a given network resource, not the *traffic load*.

Furthermore, Kirkby, explains at column 10, lines 13-16, “The WtP values stated by each user are communicated to the network manager 50 which calculates the demands for each resource and allocates bandwidth to each user accordingly,” and at column 12, lines 12-13, “Bandwidth allocated to the CPE 70 is based on the CPE's stated WtP value.” These passages reveal that in Kirkby network bandwidth allocation is done based on customers' preset, *stated willingness to pay* values, ***and not based on the actual, prevailing traffic loads***. Accordingly, Kirkby does not disclose controlling connection capacities based on the traffic loads of the connections, as recited in claim 1.

This claimed capability to control connection capacities based on their traffic load variations is greatly valuable in several network applications, including wholesale network contracts, operator inter-exchange applications and corporate inter-site networks where all network access capacity is paid for with a flat fee — i.e., where the revenue (cost) of the network contract is pre-determined — and where the main objective of connection capacity control is to maximize the utility of the network i.e. its traffic throughput. The claimed capability to optimize the connection capacities continuously based on their traffic load variations enables achieving this objective.

“each node [..] providing a connection of variable capacity to the other nodes”

Claim 1 also recites that connections between the nodes in the network system have variable capacities. Kirkby does not disclose a network system with variable capacity connections.

As explained in the previous section, network bandwidth allocation in Kirkby is done based on stated, preset willingness to pay (WtP) values provided by users ahead of time to the network controller, instead of based on the traffic loads variations. In addition to this, Kirkby further assumes that the WtP values (demands) will be *stable*, rather than *variable*, and that the network capacity distribution done based on WtPs will settle in a *steady* (i.e. *non-variable*) state of network bandwidth allocation. Kirkby explains this, for example, at column 13, lines 46-48, “The throughput traffic through a node determines how much traffic can be inserted onto the network at any instant of time. The above process is done at every instant of time and this leads to the proportionally fair traffic distribution of the network once the traffic has settled down”; at column 14, lines 52-53, “the algorithm converges to the proportionally fair distribution in the steady state”; and at column 15, lines 21-22, “the ring will rapidly converge on the exact proportionally fair solution for stable input demands.” Because Kirkby’s network connections will have *steady* (i.e. *constant*) rather than *variable* capacities, Kirkby does not disclose variable capacity connections as recited in claim 1.

The capability per claim 1 to control the variable capacities of connections dynamically based on their traffic loads is important in many applications for network traffic throughput maximization. This capability is valuable to the network operator, as the utility of the network of a given capacity (cost) is typically directly related to the throughput of traffic it is able to deliver between the sites interconnected by the network.

For at least the foregoing reasons, Kirkby does not anticipate claim 1, or any of claims 3-8, which depend from claim 1. Because the obviousness rejection of claim 2 expressly applied Kirkby in the same way that the anticipation rejection of claim 1 applied Kirkby, claim 2 is patentable over the combination of Kirkby and Dai for the same reasons set forth above.

Based on the foregoing, the application is in condition for allowance of all claims, and a Notice of Allowance is respectfully requested. If the examiner believes for any reason direct contact would help advance the prosecution of this case to allowance, the examiner is encouraged to telephone the undersigned at the number given below.

Respectfully submitted,
MARK SANDSTROM

Dated: June 28, 2006

By: /Robert A. Hulse/

Robert A. Hulse, Reg. No. 48,473
Attorney for Applicant
Fenwick & West LLP
801 California Street
Mountain View, CA 94041
Tel.: (415) 875-2444
Fax: (415) 281-1350